

40th ANNIVERSARY YEAR, 1965-2005

Programme for members' meetings 2004 – 2005

Tuesday 7 September 2004
Late afternoon

Computer Audit Basics 2: Auditing the Infrastructure and Operations

16:00 for 16:30
KPMG

Thursday 7 October 2004
Full day

Regulatory issues affecting IT in the Financial Industry

10:00 to 16:00
Old Sessions House

Tuesday 16 November 2004
Full day

Networks Attacks – quantifying and dealing with future threats

10:00 to 16:00
Chartered Accountants Hall

Tuesday 18 January 2005
Late afternoon

Database Security

16:00 for 16:30
BCS

Tuesday 15 March 2005
Full day

IT Governance

10:00 to 16:00
BCS – The Davidson Building,
5 Southampton Street,
London WC2 7HA

Tuesday 17 May 2005
Late afternoon

**Computer Audit Basics 3: CAATS
Preceded by IRMA AGM**

16:00 for 16:30
BCS – The Davidson Building,
5 Southampton Street,
London WC2 7HA

AGM precedes the meeting

Please note that these are provisional details and are subject to change.

The late afternoon meetings are free of charge to members.

For full day briefings a modest, very competitive charge is made to cover both lunch and a full printed delegate's pack.

For venue map see back cover.

Contents of the Journal

Technical Briefings		Front Cover
Editorial	John Mitchell	3
Chairman's Corner	Alex Brewer	4
The Down Under Column	Bob Ashton	5
IRMA Venues Map		5
Measuring the effectiveness of an internal control system	Dr. David Brewer and William List	6
Membership Application		25
Management Committee		27
Advertising in the Journal		28

GUIDELINES FOR POTENTIAL AUTHORS

The *Journal* publishes various types of article.

Refereed articles are academic in nature and reflect the Group's links with the BCS, which is a learned institute governed by the rules of the Privy Council. Articles of this nature will be reviewed by our academic editor prior to publication and may undergo several iterations before publication. Lengthy dissertations may be serialised.

Technical articles on any IS audit, security, or control issue are welcome. Articles of this nature will be reviewed by the editor and will usually receive minimal suggestions for change prior to publication. News and comment articles, dealing with areas of topical interest, will generally be accepted as provided, with the proviso of being edited for brevity. Book and product reviews should be discussed with the appropriate member of the editorial panel prior to submission. All submissions should be by e-mail and in Microsoft Word, Word-Pro, or ASCII format. Electronic submission is preferred.

Submissions should be accompanied by a short biography of the author(s) and a good quality electronic digital image.

Submission Deadlines

Spring Edition	7th February	Autumn Edition	7th August
Summer Edition	7th May	Winter Edition	7th November

The views expressed in the Journal are not necessarily shared by IRMA.
Articles are published without responsibility on the part of the publishers or authors for loss occasioned in any person acting, or refraining from acting as a result of any view expressed therein.

Editorial Panel

Editor

John Mitchell

LHS Business Control
Tel: 01707 851454
Fax: 01707 851455
Email: john@lhscontrol.com

Academic Editor

David Chadwick

Greenwich University
Tel: 020 8331 8509
Fax: 020 8331 8665
Email: d.r.chadwick@greenwich.ac.uk

Editorial Panel

Andrew Hawker

University of Birmingham
Tel: 0121 414 6530
Email: hawkeracj@bopenworld.com

George Allan

UNITEC
Tel: +649 815 4321 x6036
Email: gallan@unitec.ac.nz

BCS Matters

Colin Thompson

British Computer Society
Tel: 01793 417417
Fax: 01793 480270
Email: cthompson@bcs.org.uk

Events Reporter

Rupert Kendrick

Tel/Fax: 01234 782810
Email: RupertKendrick@aol.com

Australian Correspondent

Bob Ashton

Wide Bay Australia Ltd
Tel: +61 7 4153 7709
bob_ashton@excite.com

The **Journal** is the official publication of the Information Risk Management & Audit Specialist Group of the British Computer Society. It is published quarterly and is free to members.

Letters to the editor are welcome as are any other contributions. Please contact the appropriate person on the editorial panel.

Editorial address:

47 Grangewood,
Potters Bar
Herts, EN6 1SL
Email: john@lhscontrol.com

Designed and set by Carliam Artwork,
Potters Bar, Herts
Printed in Great Britain by PostScript,
Tring, Herts.

Editorial

People in general tend to use the control word somewhat loosely and so do auditors. This does not mean that auditors are not people, although real people seems to think of auditors as being some bizarre sub-species slightly below the amoeba in terms of evolution. Still, it could be worse. We could be double glazing salesmen, estate agents, or politicians, who all seem to rate lower than us in the public's esteem. Enron did us no end of bad and for Arthur Andersen it was a killer, but at least they have achieved some form of redemption. Terribly sorry the jury got it wrong, but you are now absolved, reminds one of the witch test back in the middle ages. If you didn't die you were a witch and would be killed horribly anyway. If you did die, then you were not a witch, just unlucky, and were hopefully now in a better place. It was a lose, lose situation for the victim. Which brings me back to the control thing. For the witch finder general it was a win-win situation. He had everything under control because he could predict accurately the outcome of one of his trials and this is the underpinning of control – predictability. You can only control what you can measure and all a control is, is the ability to compare what you have with what you predicted. It is simply a test that is made at an appropriate time and place in the process. So not too difficult, but it is surprising how little this explanation is taught to up and coming auditors who are therefore condemned as a result to a career of muddy thinking and woolly reporting.

So if a control is simply a test against something that we can predict it is pretty simple to see how it applies in the world of IT, or anywhere else for that matter. The gender field should contain 'M' or 'F'. The program logic tests for these attributes and if a match is achieved the transaction is allowed to continue down the road. If it fails the test, then the user is informed and the transaction is held. A refund must be in the range of £50 - £500. Anything outside of this fails the predictability test and is rejected. So how do we still end up with those demands to pay £0.00, or risk prosecution? This is because a reasonableness test has not been applied as a result of poor design, which indicates that control has to be specified as a requirement and then built into the process from design onwards.

Simply because something meets the test criteria however, does not make it absolutely correct. For example, the refund could pass the range test, but still be incorrect in itself. The gender may pass the attribute test, but the boy may really be a girl. When I was young I could never understand how my mother could peer into a pram, examine the bundle of wrinkles inside and then exclaim "what a beautiful girl". It was much later that I found out about colour coding and started to apply it myself until coming across a baby dressed in yellow. My "Oh, what a beautiful it" was not well received and it was at that stage I realised one really needed control in-depth. Sometimes we can build one test on another test to reduce the likelihood of error. Gender equals 'M' and operation type equals hysterectomy do not lie easily together, for example. No matter how complicated we make the control (test) the basics remain the same in that we are testing for what we can predict. This is how test data is used and test scripts defined. We know both the question and answer, unlike the crew in the Hitchhikers Guide to the Galaxy who only knew the answer. Someone once wrote to Douglas Adams with a long monologue as to why the answer should be 29 rather than 42. I suspect this person was an auditor with control deficiency syndrome.

The reason for my theme is that in this edition there is a marvellous paper by William List and David Brewer which takes this simple concept light-years further by examining how control effectiveness can be measured. William is a previous chairman of this specialist group and was one of my early mentors in the control field. The beauty of this paper is that the authors keep the concept simple, but still manage to ask and answer all those awkward questions regarding control effectiveness measurement. This is a little like introducing the concept of $E=MC^2$ to Isaac Newton or calculus to Berewolf without the need for all the in-between stuff, but having it still making sense to them. Ultimately it all comes down to money of course, but perhaps that is the 42 of the control world. At least here we know what the question is.

Elsewhere we have a report for our current chairman Alex Brewer and a down-under column from Bob Ashton. Enjoy your summer break. Winter is not too far away.

John Mitchell



Chairman's corner and AGM report

Alex Brewer

Help!

Please consider joining and helping out with the IRMA committee. Being involved in the committee is not onerous. The meetings are brief and to the point so as not to waste people's time. I can see by the gradually increasing number of people on the first train out from Paddock Wood that people's time is increasingly scarce which is why the committee don't waste it!



One of the perks of being a committee member is gratis access to the training that occurs as well as sight of current BCS initiatives. Please get involved.

IRMA's new home

IRMA now uses the BCS Southampton Street office for all of its meetings. For the first time we have a BCS 'home', allowing the committee to focus on other issues. If you haven't been to see it then you are in for a pleasant surprise! The previous year had one or two incidents caused by not having a regular meeting place – this is now well and truly behind us and we can settle into the BCS London office. See opposite page, or our website for maps.

Finance

A normally dry subject rises to the top of the AGM for the following reason: our finances for the year are not good (down by approximately £5,000 - the accounts can be found elsewhere in this Journal) and they are prompting some changes. In particular, one change will be to move the Journals in print from paper to PDF format to save printing and distribution costs. While our reserves are suffering, they are still adequate, but now is the time to make the change. If you would like a copy of the accounts please contact admin@bcs-irma.org.

Email campaign

Apologies in advance for what will doubtless be a fairly extensive campaign to get those email addresses in, but in order to get PDF format journals and/or links to our members we will need to gather those missing email accounts. So please provide your current email address to admin@bcs-irma.org (if you haven't already). You should see a number of reminders in the coming months. We aim to cut over to electronic distribution entirely at the end of 2005.

The email addresses will provide the future backbone for passing information to members and also seeking assistance from interested parties. For example, during the last year the

BCS sought informed comment on a number of issues at short notice. The only way in which those types of deadline can be met is by using email.

Thanks...

I would like to thank all of the current committee for their hard work during the year and ask for their continued support in the coming one. I would also like to thank all of the members who have supported our meetings in the year. If you like what you see and hear, please tell a friend. If you don't, then please, please tell the committee!

Identity theft

A small paragraph of business issues for this journal. Those who follow Bruce Schneier's 'counterpane' column will have seen an article about 'Identity theft' in which he takes an interesting stand on this issue. If you don't subscribe then visit www.counterpane.com and put your name down - the emails are always thought provoking.

He asserts that identity theft is not a theft of identity as such but an attack on organisations which do not authenticate their transactions adequately. He proposes using new models of authenticating transactions that identify questionable transactions when they are created rather than trying to prevent fraudsters from opening accounts in other people's names, or identifying those questionable accounts. Identifying those accounts is often too late to recover the money in any case.

IRMA's continuing role

To take such events as identity theft to management, and assess the impact on your organisation requires the input of alert risk management professionals, and it seems to me that IRMA's skills will be in demand for some time to come.

IRMA continues to supply:

- Training, and
- A network of people.

I think the people network is often overlooked. If you have pressing issues, why not bring them to the next meeting and discuss them with the attendees, starting with the famous words 'I have a friend'? Please make use of the people network.

Watch this space or your inbox for further details. If you have not provided us with your email address yet, please email admin@bcs-irma.org, and don't forget to mention your IRMA number.

Activities for 2004/5

During the year, IRMA ran the following courses:

- Sept - Computer Audit Basics - Auditing the Infrastructure and Operations
- Nov - Network Attacks – quantifying & dealing with future threats
- Jan - Database Security
- May - Computer Audit Basics - CAATS

The day on IT governance planned for March was pulled at the last minute, but will be going ahead in the autumn.

The Down Under Column

Bob Ashton – IRMA Oceania Correspondent



IIA GTAG and Zero-Day Attacks

The Institute of Internal Auditors has published a pre-release draft of the 2nd of its Global Technology Audit Guides entitled *Change and Patch Management Controls: Critical for Organizational Success*, which is available from: www.theiia.org/index.cfm?doc_id=4706.

The Guide is written in plain English and is aimed at educating non technical auditors on information technology change management risks and controls to enable them to become more involved in the IT governance process.

Throughout its 73 pages, the guide hammers the point that unplanned and unscheduled technology changes are harmful to business and are to be avoided. While this advice is self-evidently true, it fails to address the issue of zero-day and near zero-day attacks.

A zero-day attack can be defined as an exploit, either a worm or a virus, that arrives immediately after, or even before, the

public announcement of vulnerability in a computer system, or the malicious exploitation of previously unknown vulnerabilities.

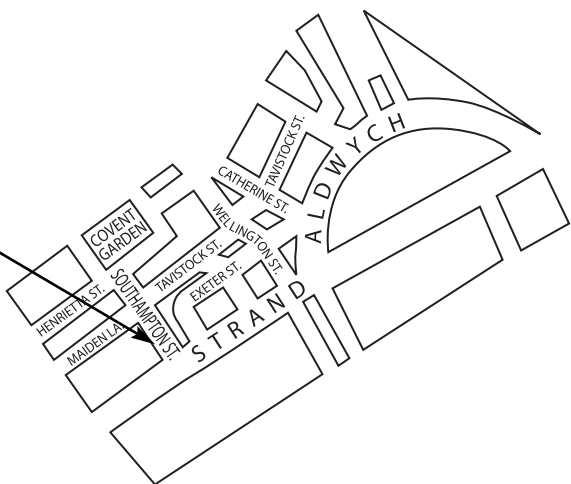
The most obvious safeguard against these attacks is to apply vendors' security patches as soon as possible after they have been released. The problem with this approach is that it does not allow any time for testing the changes before they are placed into the production environment. The Guide provides a full explanation of the risks associated such a practice. If anyone is aware of a solution to this conundrum I would be very interested to know what it is, as I am sure, would be the authors of the draft Guide.

The final publication of the Guide is planned by the IIA for July. Its utility would be much improved if the repetitive nature of some of its advice is reduced and the issue of zero-day attacks addressed.

PLEASE NOTE THE EMAIL ADDRESS FOR
IRMA ADMIN
IS:
admin@bcs-irma.org

Venue for Full Day Briefings

BCS, The Davidson Building,
5 Southampton Street,
London WC2 7HA



Measuring the effectiveness of an internal control system

by Dr. David Brewer and William List, CA, Hon FBCS



The objective of this paper is to propose a methodology by which management can measure the effectiveness of the organisation's Internal Control System (ICS). In addition the paper proposes a methodology for recording Risk Treatment Plans (RTPs), which improve the communication between risk specialists and senior management. This methodology incorporates our concepts for classifications of ICSs.

The ICS is the way in which the management deploys the organisation's resources to achieve the organisation's objectives.

The ICS exists in two basic parts:

- Procedures to perform the work necessary to conduct the organisation's business. These are called operational procedures.
- Procedures to ensure that the business is conducted as expected. These are called controls.

It is this second part of the ICS which this paper examines.

All organisations have an ICS. In large organisations it is formalised; in the very small organisations it is often implemented by the boss being involved everywhere. Most organisations are somewhere in between these two extremes.

It is axiomatic that things will go wrong - people do not always perform as expected, great new products do not sell as well as expected, criminals attack the organisation, acts of God occur, etc. This has always been the case. The conundrum facing management is to decide how much resource to deploy to create just sufficient controls to limit the possibility of bad events occurring and to limit the damage when they do occur.

When an organisation outgrows the ability of the boss to supervise everything, management have sought to resolve the conundrum by applying (a series of) risk assessments. In these assessments the probable events are identified and appropriate actions to limit damage are determined.

The question "Is this an optimum deployment?" still remains whatever controls are in place and however the need for them has been identified. The methodology we propose seeks to assist management in answering the question. It allows management to determine by direct measurement whether or not their actual ICS is achieving the objectives they want, irrespective of what else is happening in the world. In other words, the measurement is neither conditional on the frequency or other characteristics of events nor how damaging the resulting impacts might be. It allows management to measure improvements in the ICS and to tune it for overall cost-effectiveness.

In summary, we propose to measure the operational effectiveness of the control part of the ICS using various time metrics. In particular we propose to determine for each event the times relative to the time at which the event occurred (which we describe in our model as T_E):

- The time of detection (TD if detected by the ICS, or if detected by some other means TM , e.g. reported in a newspaper)
- The time that the damage caused by the event is fixed (TF), should it be possible and appropriate to fix it, or otherwise resolve the problem
- The time limit after which (TW), if the damage is not fixed, some impact penalty IP (whether financial or otherwise) is incurred.

We use the time measure because it is independent of the volume of events (which are totally variable given the threat environment) and independent of the value of events (which is random). It allows us to classify the controls as belonging to one of seven classes. We use these to determine the operational effectiveness of the ICS, which for convenience we express as belonging to one of five categories. We also use the time measures, coupled with frequency, to measure improvement in the ICS. Finally, in order to optimise the cost-effectiveness of the ICS, we introduce a set of financial metrics (or substitute metrics if financial measurement is inappropriate):

- The costs of normal operations - performing the work to achieve the business objectives (which we describe in our model as cost of doing business - CBA)
- The costs of the controls of whatever form - access control, buildings insurance, business continuity planning, IT recovery procedures, etc. (which we describe in our model as cost of the ICS - $CICS$)
- The financial impact of any events that do occur (which we describe in our model as the impact penalty IP)
- The costs of fixing or otherwise resolving the damage caused by the event (which we describe in our model as the cost of fixing the event CF).

Having optimised the operational effectiveness of the ICS, a set of inequalities using the financial metrics then allows us to tune the ICS for cost-effectiveness.

Note that those procedures which are created to facilitate recovery from an event or to minimise the impact of an event are described in this paper as a Business Continuity Plan (BCP).

In practice, an ICS addresses many different types of event, and the optimum controls for each one could fall into any one of the seven different categories. Thus a real ICS may have controls belonging to each and every category. We therefore propose a methodology for choosing the optimum controls for an ICS that must address a wide variety of different events and impacts.

The remainder of the paper is divided as follows:

- The next section presents the background to this paper
- We then recount some true stories and anecdotes that provide a foundation to our theory
- We next present the fundamental model

- We define the control classes and categories of ICS
- We then describe how to measure/monitor operational effectiveness and improvements, and tune the ICS for cost-effectiveness, with the aid of some worked examples
- We then present our methodology for generating Risk treatment Plans (RTPs)
- Finally we present our conclusions.

BACKGROUND

The Need for Control

Ever since organisations expanded beyond the control of the “boss” there has been a need for controls to regulate their activities. For example the profession of accountancy/audit grew out of the need for owners to check on their factors/agents overseas in the 19th century. As private companies expanded and brought in outside shareholders (joint stock companies) the need to regulate the behaviour of those running the companies grew and the first set of legislation governing companies was passed in the early 20th century.

Since the Second World War there has been very substantial change; the development of IT, the expansion of cheap communications (both travel and telephones) across the world etc. These new facilities have been harnessed by commerce to create world wide organisations that can be operated from one point on the globe. The need therefore to update the legal framework for the conduct of commerce (and governments, charities etc) was recognised and a large volume of laws and regulations now exist in most countries specifying standards of conduct and controls that must be complied with by organisations.

Many of the new laws are a result of scandals where it was perceived that the investing public (directly or through co-operative investments) were being “ripped off” by the inappropriate conduct of senior executives. One only has to consider the South Sea Bubble, Kruger, Salad Oil company, Equity funding, Polly Peck, Maxwell Pensions, Enron, WorldCom to name but a few to realise the potential for mischief has existed over the centuries and no doubt still exists today.

Corporate Governance

In addition a perception that the public in general, and minorities in particular require protection from the large organisations has resulted in many laws and regulations governing the conduct of organisations in relation to their employees and the public. These cover anti discrimination, privacy protection, product quality etc.

The result is that organisations require an ever more sophisticated system to ensure compliance with the laws and regulations.

In the UK the main documents covering corporate governance are the series of reports culminating in the Turnbull report (and now Higgs) which dealt with the conduct in the board rooms of UK organisations. These now are read in the context of the OECD recommendations on Corporate Governance. In the US in response to the recent scandals there is an act (Sarbanes-Oxley) that requires inter alia executives to take personal responsibility for the published material from companies.

In this paper concerning Internal Control we are concerned about the processes necessary to implement the organisation’s

mission, including compliance with the laws and regulations, and not with the details of those requirements in themselves nor specifically the Corporate Governance issues surrounding effective disclosure, fairness between stakeholders and executive remuneration.

Operational Risk

In particular we are concerned with the processes to limit operational risk within an organisation. At present the financial services regulators world wide are seeking to change the processes within the regulated organisations to accord with the Bank of International Settlement’s (BIS) requirements set out in BASEL 2. National regulators and BIS are issuing guidance on the implementation to regulated organisations.

The Need for Risk Assessment

Behind the regulatory initiatives there are a number of international standards, which affect the processes within an organisation. The three main standards today are ISO 9001 (and derivatives), ISO 14001 (and derivatives) and ISO/IEC 17799/BS 7799 Part 2 (and derivatives). ISO 9001 addresses the controls to achieve quality in products and processes. ISO 14001 addresses the controls to protect the environment. ISO/IEC 17799 addresses the processes for information security within an organisation and BS7799-2 provides the mechanisms for the management system.

The Treatment of Risk

All the regulations and standards expect organisations to establish effective controls on the basis of a risk assessment. The results of a risk assessment can be categorised as:

- Risks which require to be guarded against (i.e., the applicable risks in the Audit Practice Board Guidance)
- Risks which are either of low impact or low probability of occurrence where no specific controls are required. In the case of the very high impact and low frequency organisations often include some preplanning for an occurrence, for example business continuity planning etc. In other cases the risk may simply be deemed to be acceptable or avoidable.
- Risk where it is appropriate to transfer the (financial) implications to another organisation for example insurance, goods on consignment etc. To effectively transfer the risk it is often necessary for organisations to implement associated controls, for example to ensure compliance with the requirements of an insurance policy and to address non-financial impacts, such as the availability of office space.

Types of Control

We assert, for the purpose of explaining our theory, that a risk materialises on the occurrence of an *event*, the consequences of the event being the damage caused by the adverse *impact* (and recovery from that impact). There are three classes of controls:

- Preventive - which seek to ensure the impact never materialises. This type of control either prevents the event from occurring or affecting the organisation, or detects the event as it happens and prevents any further activity that may lead to an impact.
- Detective - which identify when some event, or events have occurred that could lead to a materialisation of the impact, and invoke appropriate actions to arrest (or mitigate) the situation.

- Reactive - which identify the impact has occurred and invoke appropriate actions to recover (or mitigate) the situation.

Certain events will not usually be able to be detected by an organisation's Internal Control System (ICS). For example, a terrorist alert requiring closure of the office will be notified by the authorities. Other events will be detected by the stakeholders - customers, suppliers, shareholders, employees etc who make complaint to the organisation when they perceive that things are wrong (perhaps incorrectly!). ICSs should therefore include processes for handling complaints fully - including identification of the cause if there was error on the part of the organisation.

Our Objectives

The problem facing the senior management with regard to the controls can be expressed as the following questions:

- Do the controls work (including are they performed correctly)?
- Are they cost effective?
- Do we have sufficient (neither too many or too few)?

Organisations monitor their controls in two main ways:

- Investigating incidents (i.e., events and impacts) and making amendments to controls as appropriate
- Conducting formal or informal audits.

Both these methods tend towards creating more controls than the minimum necessary. Reaction to incidents may be "knee jerk" and "over the top". Auditors often rightly identify problems in a control structure and suggest additional controls to fill the gaps, as they see them.

Our methodology seeks to create an objective set of measures to assist management to judge the cost effectiveness of the controls in this ever more regulated world.

SOME TRUE STORIES

"A funny thing happened to me on the way to the theatre..." This timeless phrase reminds us that it is always worth recounting some true stories and anecdotes at the outset of a serious activity. It enables us to impart some of our experiences that led us in some way to the conclusions that we have drawn.

There are six such stories. They concern:

- What happened to the authors on Flight BA122
- Why we have chip and PIN
- Tales of the unexpected
- When acceptable risk becomes unacceptable
- Over-reliance on technology
- Software intensive projects.

At the end of each story we make some observations, which we summarise after recounting all the stories.

Flight BA122

On Saturday 22nd November 2003 we relaxed back in our seats aboard BA 122 from Mauritius to London Heathrow fully expectant of the usually smooth take off as we rocketed down the runway at 190km/hr. Imagine our total shock and horror as the Captain forcefully applied the brakes, just at that point where we expected to be airborne. We must admit, stopping 370 tons of aircraft at that giddy speed in less than 7 seconds was

impressive, even if it was an experience we would all prefer never to have to repeat. As the plane shuddered to a halt – and indeed the vibration is horrendous – we were glad that all the baggage had been properly stowed and our belts were tightly fastened. Were we about to hit something? Had something fallen off? Was this the end? No. The plane stopped and all was quiet.



View from the veranda of Sugar Beach

The Captain explained the problem. It appeared that there was an engine fault.

The plane taxied back to the safety of the apron well away from the terminal building. Some one hour later, after rolls and water had been served to the passengers, the passengers were told that the flight would not proceed that day. We were at the back of the aircraft and were held there because another passenger had to be taken to hospital. Eventually we were disembarked, passed back through immigration and customs where we waited for about another hour to be allocated a room to stay the night. We were then put into an un-air-conditioned bus where the luggage was stacked onto the back seats and down the aisle. The bus took us to our hotel – Sugar Beach, about an hour's drive from the airport. Some people complained vehemently that the bus had no air-conditioning – we were just glad to be alive and thankful that BA knew what to do.

The hotel greeted us with a welcoming smile, a refreshing drink and checked us in without fuss. We had the afternoon to ourselves. England had just won the World Rugby Cup and it was gorgeously hot and sunny. We could relax for a few hours despite, as we had just been informed, that we would be woken up at the ungodly hour of 4am to get back to the airport. We were also told that BA would not pay for alcoholic drinks. We understood that, settled in and had some lunch, it being about 2.30pm by this time.

That night we settled our hotel bill after much discussion as to what was and was not to be paid for by the passengers - it transpired that lunch was not being paid by BA.

The following morning there were no calls to wake us up. Those of us who got up early were not immediately told that the flight was further delayed although the hotel must have known at that time, as otherwise they would have woken us up. Some passengers got a great fright on waking up about 6am thinking they had missed their plane. We were told to wait about for differing periods of time until about 11am when we were called back to the buses to take us back to the airport. Then we discovered that we had to pay for everything we had taken from the minibars – soft drinks and one passenger even had to pay £3 for a tea bag! Again the luggage was piled up at the back of the bus.



Waiting for the bus at 4am!

At the airport we again queued for the X-ray machine, the check-in and immigration. Amused that we now had two exit stamps in our passports we waited and waited. It was then announced that the Captain had performed his pre-flight checks and had grounded the

aircraft as the same cockpit warning lights were lit. This was indeed an unexpected surprise, as rumours had it that an engineer had been flown from London to certify the repairs and they had passed! We gathered around the gate, however many hundreds of passengers we were, to hear the Captain address us personally. We wanted to shake his hand and thank him for putting our safety first, but others were clearly unhappy. "Are we dead?" we heard someone say. "No", came a comforting reply "we practice until we all get it right, then we take off". The Captain spoke reassuringly through his megaphone. He explained the situation with the engine and informed us that there were two alternatives. One to fly out another aircraft from South Africa, the second to stay a further night in Mauritius while a new aircraft was flown out from the UK. BA provided us with a voucher for some food.

We never heard the announcement about which option had been chosen but as others on the plane had left the waiting area, we gathered that it was the second alternative. We went back through immigration and again queued to get our allocated hotel room. We were bussed to the same hotel and met there by BA staff. Three buses were required to get us there, and we arrived at sundown.

We were told that we would be woken at the even more unearthly hour of 3am to get back to the airport. Again BA did not pay for alcohol but this time there were no charges for items from the minibar.

In the morning we were woken up, had breakfast, and then waited for some 1.5 hours for the buses to come. Why, if it had required three buses to transport us to the hotel, were only two buses sent to take us back? There were insufficient places on the two busses to take all the passengers and their luggage! So after considerable muddle and much anxiety the buses left (others then got taxis) and we again went through the checking procedures at the airport and waited in the departure lounge. Indeed the comment that we practice until we all get it right seemed rather pertinent. The check-in procedure was clearly faster, everyone knew exactly what to do and we now had three exit stamps in our passports to prove it.



Discovering that 3 into 2 does not go!

Now, of course, being Monday there were two flights with the same BA 122 number and the people waiting for the Monday flight got muddled up with the people for the delayed Saturday flight. The tannoy call to board mentioned the delayed flight but gave a date not the day, which was unclear to those who had not been delayed. The different coloured boarding cards allowed the people to be sorted out but it caused delay in the line to board.

The Captain proudly announced on the tannoy that all was well and we would board in 10 minutes so we lined up. Some 1 hour later we eventually boarded - no one seemed to know what the hold up had been.

When we got home we discovered that the BA London information had been telling the people collecting us that the flight was OK on the Sunday, which it was not and on the Monday that there was only one flight! Fortunately we had a mobile phone so kept our families informed!

BA had given many passengers a form to complete. We had to ask for ours on the aircraft. We completed it and duly sent it off. BA replied, apologising and offering us a complimentary round-trip ticket to any destination of our choice. The letter formally confirmed the delay for insurance purposes and acknowledged that the passengers should have been kept better informed, that matter having been already taken up with Senior Management in order to avoid a recurrence.

Observations

In business terms the event was "One of our aircraft has broken down in the Indian Ocean"; the impacts, *iter alia*, being "air crash", "increased costs" and "customer dissatisfaction". BA's concern for aircraft safety is undisputed, and the steps taken to avoid the "air crash" impact clearly took priority over every thing else. Having done so, however, the poor communications and apparent succession of short-term decisions gave an appearance (at least at the time) of minimising "increased costs" over "customer dissatisfaction". The apparent short-term decisions were:

- Fix the engine locally, fly out an engineer from the UK in parallel to certify the repair. The cost to fix is then the cost to put up everyone for one night plus repair etc.
- If that does not work, fly out a plane from South Africa. The cost to fix is then the cost to put up everyone for one night plus repair plus cost of plane from SA and its consequential costs etc
- If that does not work, fly out a plane from the UK. The cost to fix is then the cost to put up everyone for two nights plus repair plus cost of plane from SA and its consequential costs etc.

Had it been decided to fly out an aircraft from the UK immediately, and to keep everyone informed with a single plan that is guaranteed, would it had given greater customer satisfaction? As the people on the receiving end, we think "yes". Would it have cost more? As things turned out, probably not, probably considerably less! Thus, as the story unfolds, we see an apparent *balancing act* between the costs of doing things to mitigate/fix the problem and the financial ramifications of the resulting impacts. Part of this balancing act is getting the *priority ordering* of the impacts right.

We invited BA to read the story above before publication. They correctly pointed out that the story recounts our experiences of what happened. It does not necessarily reflect what BA intended to happen. We must remember that for much of the time we were in the hands of BA's agents, rather than BA itself, and the agents may or may not have carried out BA's instructions in the way BA had intended. We do not know how much of the groundside disorganisation was due to the Airport and the handling agent and beyond BA's control. Perhaps BA paid for the bus that never turned up, as well as footing the bill for all the taxis called to replace it. Perhaps BA asked for air-conditioned buses. If BA was making risk management decisions in London based on "safety first, customer second, cost third", perhaps, like us, it had inadequate information. A free air ticket to anywhere in the world is a pretty magnificent gesture of compensation - but not the best way to achieve "safety first, customer second, cost third".

In this example, the combination of on-board electronics and pilot competence clearly illustrates the ICS was able to detect the initial event in sufficient time for something sensible to be done about it. It also shows that in cases such as this very fast reactions are required. Subsequently we find:

- The initial plan to deal with a jumbo full of people stranded at the airport worked well.
- The transport of the people was as best it could be in the circumstances.
- The communication by the hotel on the second morning was poor and the payment arrangements were a muddle.
- The communication at the airport on the Sunday and Monday was poor and disorganised.

Overall the plan, which started so well, seemed to fall apart the longer the delay in the flight took and the more different people were involved. We deduce that the ICS's ability to cope with the consequences of further complications after the initial event was poor and may have involved decisions made without full information or without full consideration of the overall impact.

Chip and PIN

Credit card fraud has existed for as long as credit cards have existed. The payment associations (VISA, MasterCard etc) are pretty much on the ball and use quite sophisticated techniques to track down the culprits whilst protecting their members' customers.

Until recently, however, making suggestions on how to improve security pretty much fell on deaf ears. To the mind of a security practitioner, the amount of money that was regularly lost due to fraud seemed infinitely large compared to the cost of the information security services that were being offered to combat the problem. What seemed stranger was the argument that the loss was small fry compared to the billions of dollars that were being transacted every day. In other words, it was an acceptable risk. However, with the widespread introduction of "chip and PIN", it would appear that the risk is no longer acceptable.

Chip and PIN means using a smart card with cardholder authentication provided through a traditional 4-digit PIN. The GlobalPlatform technology serves as a good, well thought out example in the context of dynamically reconfigurable smart cards. Compared to a magnetic stripe card, the smart card is significantly harder to clone and persuade to divulge its secrets (e.g. the PIN). GlobalPlatform cards are able to defend themselves against attack and can communicate with the Card Issuer. Thus:

- Individual applications can be blocked, e.g. for every cardholder, if a security weakness is discovered in that application. Subsequently, the vulnerable applications can be deleted and replaced by a new version that does not exhibit that vulnerability.
- As is the case now with magnetic stripe cards, an individual card can be blocked, e.g. if reported lost or stolen, or suspected as such.

Thus the objective of chip and PIN is to reduce the number of attempted fraudulent transactions, by introducing a more reliable cardholder authentication mechanism, that is also extremely difficult to tamper with.

Observations

By itself chip and PIN will not, and cannot, reduce the set of attempted fraudulent transactions to zero. It will not stop the thief who guesses the PIN, or found it conveniently written down in the gentleman's wallet. It will not stop the genuine cardholder from spending more than the Card Issuer is willing

to lend them. Other controls, which already exist such as authorisation limits, are necessary to do that. What is does do, however, is (a) decrease the time between the event (attempted unauthorised use) and its detection; (b) increase the reliability of that detection.

In the event that someone forges the cardholder's signature sufficiently well for the shop keeper not to notice, the point at which the unauthorised use of the card is discovered could be days after the transaction has taken place. The goal of chip and PIN is to render such detection virtually instantaneous. Thus the decrease the time between the event and its detection afforded by chip and PIN is significant. It *detects the event so fast* that all subsequent activity, which would otherwise lead to the occurrence of some adverse impact, is prevented. It is therefore a preventive control. In contrast, the controls that traditionally spot fraudulent activity detect the event too late, the impact having already occurred.

The cost of rolling out chip and PIN is not insignificant, but so is the cost of credit card fraud. The introduction of chip and PIN shows that the *balance* between the cost of control and the cost of impact has shifted in favour of greater control.

Tales of the Unexpected

An organisation had built a brand new European Headquarters which conformed to the best practice for construction and Health and Safety regulations. The building was equipped with sprinklers and extinguishers as well as being constructed with fire proof material. Clearly these matters form part of the ICS of the organisation in that they were costs incurred to guard against the unlikely eventuality of a fire, even though most were compulsory to comply with regulations. In addition, following previous experiences with fires the organisation had in place a tested recovery system for the head office IT systems and applications and procedures for dealing with personnel issues, the press, loss adjusters etc in case of disasters. In effect they had in place an ICS including BCP, some of which was in place and some tested but only activated as required.

Unfortunately there was a small fire in one wing of the building and the fire procedures were invoked including calling the fire brigade. During the course of setting up the fire fighting equipment the wrong water valves were used and the sprinkler system was inadvertently turned off; the result was that the fire spread rapidly in the roof space to the whole building. Now there was a disaster, not merely an inconvenience because the Head office had to be relocated urgently, which was not part of any extant plan!

Observation

Controls *do not always work* as intended, and in this case with potentially catastrophic consequences.

Acceptable Risk?

The Audit Practices Board (APB) presents an interesting example of acceptable risk.

Basically, the example concerns a small advertising agency. Small adverts are placed for cash and the company accepts the risk that £5,000 worth of cash transactions may be lost per annum, for whatever reason. The APB example argues that the cost of the controls necessary to assure each transaction would be disproportionate to the value of the transactions. The problem of such losses is subsequently ignored.

Our question is “How does the company know when the loss becomes £5,001?” Surely, that ought to be an unacceptable risk!

Observations

What the APB example fails to argue concerns when this acceptable risk becomes an unacceptable risk, i.e. when the loss becomes £5,001. First, of course, you need a way to determine when it does. A reconciliation, each month, of the cash received versus the advertisements would serve this purpose. It would highlight the total loss, albeit being unable to identify the particular transactions concerned. However, it is the total that we are interested in at this stage of control.

If the reconciliation, performed at the end of month 11, shows that the loss is £4,580 then the loss remains acceptable (as it is just on target to come under £5,000) and the company can be satisfied with its decisions. If the same loss in reported at the end of month 1, then the company ought to be concerned that its acceptable loss is in danger of becoming an unacceptable loss in month 2, and ought therefore to take action accordingly. Once again, it is necessary for the ICS to detect the event (in this case the metamorphosis of acceptable to unacceptable risk) in *sufficient time* for something to be done about it.

Over-reliance on Technology

At a meeting, our client’s IT manager asked why his networks had just been the victim of a well known virus. We asked some questions and sent him off to find the answers. During his absence, a colleague remarked that for some time his laptop had been reporting that its anti-virus library was not up-to-date. Others quickly reported the same. The IT manager reported back. Anti-virus library upgrades were being received in a timely manner by the server but due to a software problem they were not being distributed to any other computer on the network. The software had stopped functioning three months ago!

With another client, we asked some questions to determine whether the anti-virus libraries were up-to-date. They were, save for all the directors’ laptops. Further investigation revealed that they were scheduled for a regular update every day at 05:30. No director had ever docked their laptop at that the unearthly hour in the morning. Their libraries were two years out of date! We asked about their new web-surfing controls. The QA manager, a railway model hobbyist, proudly announced that it prevented him access to his hobby sites, and having been denied once he had never tried again. We asked him to try once more, and guess what - he had access. The software had stopped working.

Observation

These stories remind us that controls *do not always work* as intended and from time to time they fail, but does anyone ever check!

Software Intensive Projects

We were always taught as young computer programmers of the urgency of discovering your mistakes early on in the development lifecycle. A design error found at the design stage is usually quicker and less expensive to fix than if it is discovered by the client when the system is operational! - but that depends on who is paying. For example, much of the UK government procurement for software intensive projects prior to the early 90’s was performed on a time and materials basis, and quite often overran with a corresponding escalation of costs, which the client paid for. The joke at the time concerned a conversation

between a small boy and a genie. The boy wanted to get rich. The genie replied “I’ll make you a sultan”. The boy asked to be made richer, and the genie would offer a more powerful position. Following some iteration the boy insisted that he wanted to get really, really rich, whereupon the genie would reply “I’m sorry, but there are no vacant positions for defence contractor”. Thus, the regime of time and materials contracts for many government procurements came to an end.

The initial shift was to fixed price, and in many cases, even for small contracts (<£100K), there was a requirement for a risk analysis. Thus the client:

- by insisting on a fixed price, aimed to pay the same amount irrespective of whether the contractor made a mistake or not.
- by asking for a risk analysis, presumably aimed to gain some feeling for the effectiveness of the ICS and to assure himself that there were sufficient controls in place to guard against non-delivery.

This dramatic shift of risk ownership from client to contractor met with some problems, not least what to do if the error was made by the client. This resulted in other procurement strategies, such as the Private Finance Initiative, where the risk is shared.

Observations

Quality controls are equally part of the ICS as are financial, security and environmental controls. The thrust of good software engineering techniques is generally towards *detecting errors early enough* in the development lifecycle to do something, without disproportionate expenditure of resource, to correct them. Even so, there is a cost of which has to be *balanced* against the cost of failure.

Summary

Our observations in respect of each of these stories have much in common. They are:

- Without loss of generality, an ICS must detect the event in *sufficient time* for something to be done about it. (See BA122, Chip and PIN, Acceptable risk and Software Intensive Projects.)
- Controls, irrespective of whether they are preventive, detective or reactive, *do not always work*. (See Tales of the unexpected and Over-reliance on Technology.)
- Controls cost money. So can an impact. In practice deigning the most effective ICS is likely to be a *balancing act* between the two. The *priority order* in which impacts are dealt with may also be important. (See BA122, Chip and PIN and Software Intensive Projects.)

Of these the most significant is that the time taken to detect the event must be fast enough for something to be done to prevent or otherwise mitigate the ensuing impacts. Referring back to our opening remarks on corporate scandals (see page 2), we ask whether there were any controls in place to detect the initiating event(s). If so, then clearly they were unable to prevent the consequent actions that led to such disastrous impacts, but could they have done so? If the answer is truly no, then could they have detected any of the events in sufficient time for someone to have done something to arrest the situation? Perhaps they did, but no one took any notice, or, as we would like to believe, failed to recognise the significance. Armed with an understanding of our fundamental theory (described next) and some tricks of the trade, such as event-impact analysis (see later), perhaps they would.

FUNDAMENTAL MODEL

In this section we introduce our Fundamental Model. Let us start by supposing that an organisation carries out a range of business activities. Let the cost of such activity be C_{BA} . Cost may be expressed in terms of money and/or resources (e.g. volunteer work). It will generate some business benefit B . If the organisation is a company, then B corresponds to profit, P , and is related to the cost of the business activities through revenue R :

$$P = R - C_{BA}$$

The organisation deploys an Internal Control System (ICS). This has an associated cost, C_{ICS} , which increases the cost of doing business

$$C_{BA} + C_{ICS}$$

In the context of a company this has the effect of reducing profit, see Figure 1.

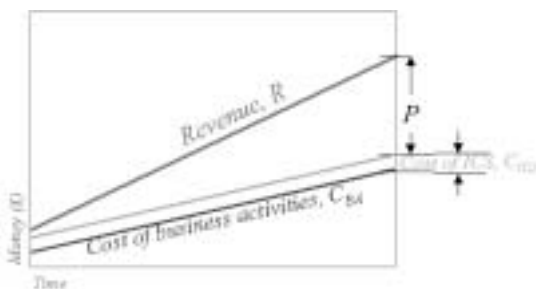


Figure 1: Impact of ICS cost on profit

Let E be a set of events: $E = \{e1, e2, e3, \dots, ej, \dots\}$.

Each event ej occurs at some time TEj and if the damage that it causes is not fixed by time TFj , where TFj is less than some time TWj (where $TWj = TFj - TEj$ is referred to as the time window), the event will cause a loss of business benefit, IPj (referred to as the impact penalty). See Figure 2.

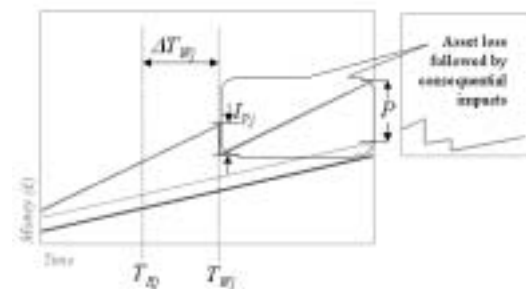


Figure 2: The onset of an impact penalty expressed in terms of financial loss

The impact penalty may take a variety of guises. For example, it could:

- arise in the form of liquidated damages or the cost of borrowing money to replace missing revenue or assets.
- correspond to reduced revenue because customers do not pay for goods or services already received or in production (e.g. as with a stage payment).
- contain hidden costs (which accumulate in C_{Fj} , see below), for example because customers demand more attention.
- be in a form that is impossible to interpret in financial terms, such as loss of life, losing the election or a court case.

Moreover, the event may also have an immediate impact on the net worth of the organisation, for example because property is destroyed or money is stolen. For simplicity, we model these asset losses as an impact penalty. As shown in the insert in Figure 2, there may also be consequential impacts, for example other customers in the future do not buy, the stock markets collapse, there is a general strike, etc.

The objective of an ICS is to control activities and detect unwanted results. An ICS is never perfect and therefore certain events will not be detected by it. Those it does detect are detected at times TDj (where $TEj < TDj$). See Figures 3 and 4.

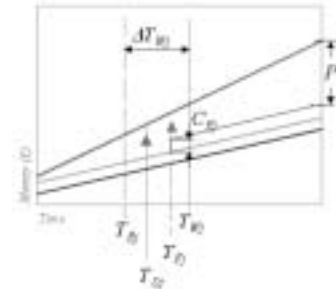


Figure 3: Detecting the event in good time to avoid the impact penalty. Impact expressed in financial terms

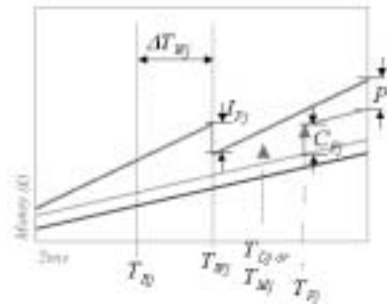


Figure 4: Detecting the event too late to do anything about it within the time window. Impact expressed in financial terms

If the ICS does not detect the event, Management is deemed to be cognisant of the event at time TMj (where $TEj < TMj$). See Figure 4.

The cost of the ICS detecting the event is included in C_{ICS} .

The cost of fixing the damage caused by the event is C_{Fj} . See Figures 3 and 4.

The damage cannot be fixed unless the associated event has been detected, i.e. $TDj < TFj$ and/or $TMj < TFj$. See Figures 3 and 4.

The impact of the event depends on when that event is detected. Specifically:

- When $TFj < TWj$ the impact is C_{Fj} . See Figure 3.
- When $TFj \geq TWj$ the impact is $C_{Fj} + IPj$. See Figure 4. Note that in this second case the time at which the event is detected TDj (or indeed TMj) may be within TWj . The problem is that the event is detected too late for anything to be done about it within the time window and consequently an impact penalty is incurred as well as the cost of fixing the damage.

The impact of the event could have a widespread effect until the situation caused by the event has been corrected; in extremis putting the organisation out of business, and/or causing widespread damage external to the organisation. In

Self-Policing Procedures

The defence is to have some other control to address failures in the first. In practice there will be a sequence of controls as illustrated in Figure 7.

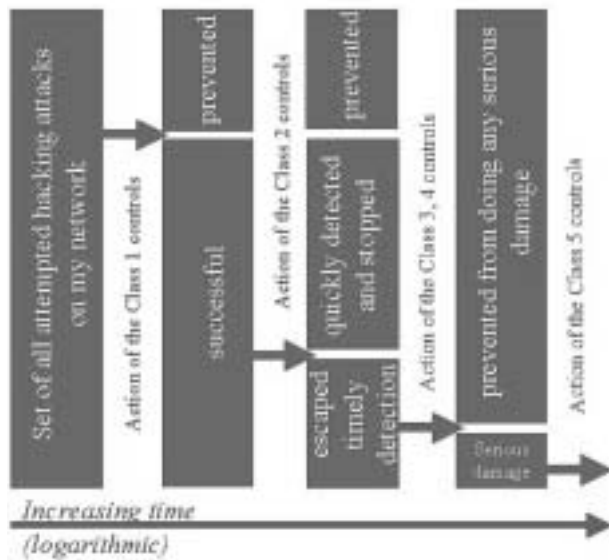


Figure 7: A sequence of controls to defend against control failures

Such a sequence is known as a self-policing procedure. It is a sequence of controls that have been constructed so that any error or failure perpetrated during execution is capable of prompt detection.

Initial detection is performed by a Class 2 control. It must be Class 2 in order to guarantee prompt detection and give sufficient time for the appropriate action to be taken before expiry of the time window.

As an example, consider a network monitoring system. When there is a failure it raises a “problem flag” and automatically sends this to the engineers responsible for fixing that type of problem. When the problem has been fixed, the engineers clear the problem flag. The engineers can falsely claim to fix the problem, but they cannot clear the alarm that raised the flag in the first place without actually fixing the problem. Think of this as a safety interlock. In addition, if the alarm is not cleared within a specified time, another problem flag is raised and sent to a higher level of management. Thus, falsely claiming to have fixed a problem or not fixing it at all does not silence the alarm but merely escalates it to a higher level of management. This is a rather well-honed example of a fail-safe self-policing procedure. Note, however, that if corrective action is never taken, the overall procedure degrades to a Class 4.

Unanticipated Events and Impacts

If there is an unanticipated event or impact it is possible, by good luck or sound judgement, that the ICS will contain something that deals with it most satisfactorily. If not, we need, almost by default, a Class 7 control to deal with it. Such a control, in some circles, is referred to as an *ad hoc* procedure.

Effectiveness Principles and Criteria

Extremes of effectiveness and ineffectiveness

Let us start by imagining what the most operational ineffective ICS might look like:

- Whatever controls it did have, if they did not work you would not find out until it was too late.
- Indeed, all the detective controls would be so slow to detect an event that the time window would always expire before the problem could be fixed.
- There would be no BCPs. When an incident happened, management would always be unprepared.

In contrast, let us imagine what the most operationally effective ICS might look like:

- Whatever controls it had, if they did not work you would find out immediately and be able to take appropriate action well within the time window. In fact all of the controls would be fail-safe self-policing procedures.
- Indeed, all the detective controls would work so fast that they would be Class 2 non-degradable. The reactive controls would all be Class 5.
- The BCPs would be so comprehensive that, when an incident did happen, management would always find that its existing Class 5 BCPs would deal with the problem entirely.

Each of these two extremes describes three principles by which we can judge the operational effectiveness of an ICS. We call them respectively robustness, speed and anticipation:

- The robustness of the ICS in the event of a control failure
- The speed at which the ICS can react to events
- The ability of the ICS to deal with the unexpected.

Some middle ground

We hope that no organisation ever has to suffer such an ineffective ICS as described above. In the categorisation below we will therefore exclude it. Likewise, we exclude the most effective ICS described above as it is too perfect. We therefore postulate some middle ground, which ideally ought to reflect good practice, and base our criteria around that. We propose for the middle ground:

- There would be some self-policing procedures, some of which may be fail-safe. [robustness].
- There would be a mixture of Class 2, 3 and even Class 4 detective controls. The Class 2 and 3 controls that were not protected by fail-safe self-policing procedures may degrade to Class 4. [speed].
- There would be at least one Class 6 BCP dealing with some catastrophe such as fire. Other incidents would be dealt with through an ad hoc procedure. [anticipation].

Above and below average

If we now think of the middle ground being some average, then we can contemplate some ICS which is below average and one that is above average. Below average, would perhaps mean that the ICS fails one of the middle ground criteria. Well below average would imply that it fails on two, but not all of them because that would describe our worst case position, which we wish to exclude. Likewise, we can consider an ICS that is above average as being one that exceeds one of the middle ground criteria. Well above average therefore exceeds two or more such criteria.

Robustness

The middle ground criterion is:

R1 - There are some self-policing procedures, some of which may be fail-safe.

A stronger criterion is:

R2 - There are some self-policing procedures, at least one of which is fail-safe.

Speed

The middle ground criterion is:

S1 - There is a mixture of Class 2, 3 and even Class 4 detective controls. The Class 2 and 3 controls that are not protected by fail-safe self-policing procedures may degrade to Class 4.

A stronger criterion is:

S2 - There is a majority of Class 2 detective controls, with possibly some Class 3 or even Class 4. The Class 2 and 3 controls that are not protected by fail-safe self-policing procedures may degrade to Class 4.

Anticipation

The middle ground criterion is:

A1 - There is at least one Class 6 BCP dealing with some catastrophe (e.g. fire). Other unexpected events incidents are dealt with through an ad hoc procedure.

A stronger criterion is:

A2 - There are a variety of BCPs (some of which may be Class 5) dealing the failure of control or some catastrophe (e.g. fire). Other unexpected events incidents are dealt with through an ad hoc procedure.

Categories of ICS

We can apply the criteria and determine the category of the ICS using a simple marking scheme. We award 3 marks for each of R1, S1 and A1 and award 1 extra mark if it is exceeded.

The resulting categorisation is:

- Well above average (AAA rating) 11 or higher
- Above average (A*) 10
- Average (A) 9
- Below average (B) 6 - 8
- Well below average (C) 4 or lower.

Example 1

To achieve a AAA rating, we need to satisfy all three criteria and surpass at least two. Thus we gain 3 marks for each criterion that is satisfied, giving $3 \times 3 = 9$, plus 1 mark for each criterion exceeded, giving $9 + 2 = 11$. If we exceed all three criteria, then the total mark is $9 + 3 = 12$, i.e. for AAA rating we need to score 11 or higher.

Example 2

To achieve a B rating, we fail one criterion, but we might exceed either or both those that we pass. We achieve 3 marks for passing a criterion and 4 if we exceed it. Thus, for a B rating, at worst we just pass two (i.e. the total mark is $3 + 3 = 6$) and at

best we exceed two (i.e. the total mark is $4 + 4 = 8$). If we pass two and exceed one, the total mark is $3 + 4 = 7$. Thus the range of marks that give a B rating are 6 - 8.

OPERATIONAL EFFECTIVENESS

Objective

Management needs to know whether or not the current ICS, i.e. the one actually in place and working now, is achieving the objectives they want, irrespective of what else is happening in the world. In other words, the measurement should not be conditional on whether or not anyone is trying to attack the organisation or defraud it, etc., how frequent those events or how damaging the resulting impacts might be. They therefore need a measure of the ICS which is direct (i.e. it is a measurement performed on the actual implementation, rather than the design) and is independent of what the world is doing. We refer to this as *operational effectiveness*. We have identified the metrics in our fundamental model. They are the time parameters:

- The time of detection (TD if detected by the ICS, or if detected by some other means TM, e.g. reported in a newspaper)
- The time that the damage caused by the event is fixed (TF), should it be possible and appropriate to fix it, or otherwise resolve the problem
- The time limit after which (TW), if the damage is not fixed, an impact penalty is incurred.

Measuring

Measurement of operational effectiveness is straightforward. It takes the form of:

- Determining the actual control class of each control, using Table 2
- Applying the criteria specified in Table 3.

R1 - There are some self-policing procedures, some of which may be fail-safe.

R2 - There some self-policing procedures, at least one of which is fail-safe.

S1 - There is a mixture of Class 2, 3 and even Class 4 detective controls. The Class 2 and 3 controls that are not protected by fail-safe self-policing procedures may degrade to Class 4.

S2 - There is a majority of Class 2 detective controls, with possibly some Class 3 or even Class 4. The Class 2 and 3 controls that are not protected by fail-safe self-policing procedures may degrade to Class 4.

A1 - There is at least one Class 6 BCP dealing with some catastrophe (e.g. fire). Other unexpected events incidents are dealt with through an ad hoc procedure.

A2 - There are a variety of BCPs (some of which may be Class 5) dealing the failure of catastrophe (e.g. fire). Other unexpected events incidents are dealt with through an ad hoc procedure.

Table 3: Summary of ICS Categories and Criteria

Award 3 marks for each of R1, S1 and A1 and award 1 extra mark if it is exceeded.

The resulting ICS categorisation is:

- Well above average (AAA rating) 11 or higher
- Above average (A*) 10
- Average (A) 9
- Below average (B) 6 - 8
- Well below average (C) 4 or lower

A Worked Example

Consider a small software company that produces bespoke software system for its clients. The company relies on an ICS that is predicated solely on program testing. In particular, there are no formal design/code reviews. There is a reliable backup system that verifies that backups are restorable and complains if they are not. However, there is no BCP covering anything outside of IT. What is the operational effectiveness of this approach?

A typical development schedule is shown in Figure 8. The “program testing” control takes effect late on in the schedule. It may start to identify problems as early on as month 6, but some problems might not be detected until month 12. If the control identifies a top level design error then the later it is detected the greater the chance that it will be too late to do anything about it before the expiry of the time window, which we will associate with the end of the development period. Thus, the “program testing” control is Class 2, potentially downgrading to Class 4. It does not tell you if it fails to find an error and therefore it is not self-policing. The backup control, however, is self-policing but is not fail-safe.



Figure 8: A typical development schedule

We have a self-policing procedure (R1 satisfied, score 3). There is nothing to suggest that there is a majority of Class 2 detective procedures, but the control under discussion is Class 2 degradable (S1 satisfied, score 3). We have a BCP (A1 satisfied, score 3). The total score is 9 and therefore we can rank the ICS as Category A.

COST EFFECTIVENESS

Objective

Operational effectiveness does not necessarily imply cost effectiveness. To determine the cost effectiveness of the ICS we need to apply other metrics, e.g. Cics, as identified in the fundamental model.

Measurement

Measurement in this case is effected by taking each control and comparing its financial parameters with those that would apply to controls that belong to other classes. There may be a variety of objectives in undertaking these measurement. Some of the most important, albeit general questions are summarised in Table 4 and discussed in the ensuing subsections. Other objectives may be very specific. Our worked example is a case in question and considers the design of an ICS to maximise the profit on a particular contract.

#	Question
1	<p>Should we be using be using a preventive control?</p> <p>Ask “Is the cost of using a preventive control less than the sum of cost-to-fix and possible impact penalties for all the events that the preventive control is designed to detect?” If the answer is yes, then there is indeed a case for using a preventive (i.e. Class 1) control.</p>
2	<p>Should we improve the efficiency of our detective controls?</p> <p><u>Upgrade from Class 4 to Class 3</u> Ask “Is the cost of the upgrade less than the average impact penalty times the number of events?” If the answer is yes, then an upgrade from a Class 4 to a Class 3 control is worthwhile.</p> <p><u>Upgrade from Class 3 to Class 2</u> Ask “Is the cost of the upgrade less than the average reduction in the cost-to-fix times the number of events?” If the answer is yes, then an upgrade from a Class 3 to a Class 2 control is worthwhile.</p>
3	<p>Should we pre-deploy our BCPs?</p> <p>Ask “Is the cost of pre-deployment over Y years minus the business benefit prior to invocation less than the reduction in impact penalty, minus the loss in business benefit, multiplied by the number of times the BCP might be invoked in that period of Y years?” If the answer is yes, then pre-deployment is worthwhile.</p>
4	<p>Should we have a BCP?</p> <p>Following consideration of the impact penalty and likelihood of occurrence, ask “Is his an acceptable risk?” If the answer is no, then you need a BCP.</p>

Table 4: Determination of cost-effectiveness

Preventive Control

Should we be using a preventive control? The answer is likely to be yes if the cost of using a preventive control is less than the sum of the cost-to-fix and possible impact penalties for all the events that the preventive control is designed to detect, i.e.

$$C_{preventive\ control} < \dots_{all\ events,\ j,\ that\ the\ preventive\ control\ detects} (Cf_j + Ip_j)$$

The cost of using a preventive control includes the cost to buy/develop, install, configure, commission, operate, train people in its use, audit its use and maintain it. We have not included in the above formula the cost of the controls that would otherwise perform the task of the preventive control as we recommend that they be retained in case of a control failure in the preventive control. An impact penalty will only occur if the corresponding existing controls are Class 4 or lower.

Note the summation over all the events that the preventive control is designed to detect. In practice, this number will be an estimate as it concerns the future. If the cost-to-fix and possible impact penalty for each event is constant (or can be considered approximately so), the inequality becomes:

$$C_{preventive\ control} < N * (Cf + Ip)_{average}$$

where N is the number of events. If this number is very small and average cost-to-fix and impact penalty is also small, then it is very likely that a preventive control will not be cost effective.

On the other hand, if either N or the average cost-to-fix and impact penalty is very large, the use of a preventive control is most likely to be a very good idea.

Detective Controls

Should we improve the efficiency of our detective controls? If the control is a Class 4 then it detects the event too late for it to be fixed within the time window. There is therefore a cost-to-fix and an impact penalty. If we convert the control to a Class 3 then there is no impact penalty. Thus, for the upgrade to be worthwhile, the cost of the upgrade must be less than the average impact penalty times the number of events, i.e.

$$C_{upgrading} < N * IP_{average}$$

We have assumed here that the cost-to-fix is about the same for the two classes of control. This is a reasonable assumption if the Class 3 allows the event to be fixed just prior to the expiry of the time window while the Class 4 fixes it immediately after. The cost-to-fix, however, may be dramatically reduced as we reduce the time taken to detect the event, thereby moving from a Class 3 to a Class 2 control. This upgrade is worthwhile if the cost-to-fix of the upgrade is less than the average reduction in the cost-to-fix times the number of events, i.e.

$$C_{upgrading} < N * (C_{fin\ Class\ 3\ case} - C_{fin\ Class\ 2\ case})_{average}$$

Pre-deployment

If we pre-deploy all or part of a BCP there will be an associated pre-deployment cost and a maintenance cost. Pre-deployment costs may include equipment purchase/lease, building purchase/hire, insurance, extra staff, training, commissioning, regular tests and practices, etc.

Prior to invoking the BCP there may also be an associated business benefit, B_{BCP} , which will offset the pre-deployment costs, $C_{pre-deploy}$. For example, a redundant IT installation, kept in a state of readiness in case the main installation fails, or is otherwise rendered unavailable, could be used for other purposes, e.g. systems development. Of course, when the BCP is invoked this benefit will be lost, hopefully for a short time but it will be lost. We need to factor this loss into our cost-effectiveness considerations. We will do this in a moment. Let that loss be $B_{BCPloss}$.

By pre-deploying the BCP we gain time. Recovery from abnormal to normal operations will be quicker, and the impact penalty will be reduced. Let that reduction be $IP_{reduction}$.

Arguably for pre-deployment to be cost-effective:

$$C_{pre-deploy} - B_{BCP} < IP_{reduction} - B_{BCPloss}$$

Invocation of the BCP ought not be a frequent event, otherwise we should be considering Class 2, 3, or 4 controls as our main line of defence. Suppose we estimate that over a period of Y years the BCP is invoked N times, our inequality then becomes:

$$-Y_{years} (C_{pre-deploy} - B_{BCP}) < N * (IP_{reduction} - B_{BCPloss})$$

BCP Need

Do we need a BCP? is what would be the impact penalty if the unthinkable was to happen? Having answered that question, the next question is "is that an acceptable risk". If it is not, then you need a BCP.

Surprisingly, in response to the second question the answer "if I am still alive, I'll just start up again" is a Class 6 control. The question is then, how much of this plan should we pre-deploy (clearly following carefully consideration and suitably refinement of this somewhat embryonic/flippant BCP!).

A Worked Example (continued)

The foregoing allows us to answer general, albeit important, questions about the effectiveness of our ICS. By way of our worked example, however, we look at something very specific and entirely in tune with what the focus of cost-effectiveness for a commercial organisation ought to be. i.e. *profit*.

Let us suppose that our small software company is being invited to bid for a new project. The contract is for a fixed price with stage payments. The development is scheduled to last one year. The final payment is due on satisfactory completion of the project. If there is an overrun, there is an impact in the form of a revenue penalty which increases with the extent of the overrun. The agreed schedule is shown in Figure 8.

As previously mentioned, the company usually relies on an ICS that is predicated solely on program testing. In particular, there are no formal design/code reviews. In costing the project, the company anticipates using a full time team of three analyst/programmers of the same grade and salary costs. Expressed in terms of some arbitrary monetary units (MU), the cost of the project is therefore 36MU plus some allowance for overhead, charged at 5MU per person per year. This gives a total of 51MU. To remain competitive the company wishes to charge the client 60MU, yielding an anticipated profit of 9MU. There is also a one-year maintenance component of 10 MU, which applies for the 12 months following client acceptance. Its purpose is to fix program bugs that manifest during the operational use of the software.

From experience, the company realises that the most likely worst case scenario is a top level design error (the event) that causes rework affecting 1/3rd of the program modules. It estimates (see Figure 9) the cost (in MU) to fix the problem as a function of the month in which the error is detected by the ICS. Past experience also indicates that all of the 10MU maintenance provision is used up, resulting in effectively no profit or loss.

Before agreeing the contract, the company considers alternative forms of ICS. It also notes that the revenue penalty is quite steep, being 5MU for any overrun, increasing thereafter by 1MU per month.

Month	1	2	3	4	5	6	7	8	9	10	11	12
Cost (in MU)	1.5	1.5	1.5	2	3	4	4.5	5	5	5	5	5

Figure 9: Cost (in MU) of fix as a function of month in which event is detected

It reasons:

Case 1

Based on past experience, the company believes that if the event occurred, at worst it would be detected in month 11 during the integration testing, but no later. According to Figure 9, this would correspond to a cost to fix of 6MU. The company also realises that, given the late detection of the event, there would be an overrun of 2 months, causing a revenue loss of 6MU, together with an overhead component of $3 * 5 / 12 = 2.5MU$. Thus, the anticipated profit margin of 9MU is at risk of being reduced to a loss of 5.5MU, should the event occur.

Case 2

The company argues that through the use of certain more sophisticated testing techniques, at an extra staff cost of one extra analyst/programmer, half time from month 5 onwards, plus a one-off software purchase of 2MU, at worst the error would be detected in month 6. The company also believes that the improved testing techniques will have a positive impact on the maintenance phase, which ought to result in a profit of 5MU, whereas usually there is none.

The increased costs of the ICS are, in this case, 5.7MU. The cost of fix, should the error occur is 1.5MU. Thus, the overall profit (inclusive of the maintenance component of the contract) would be 6.8MU should the event occur and 8.3MU should it not.

Case 3

The company recognises that an alternative approach would be to entertain design/code reviews. The company decides that this can be accomplished with the proper use of certain overhead resources on a very part-time basis in order to moderate the reviews. The project team must, however, be trained in the techniques that are to be used. The training cost will be 2MU. In the worst case, the event should be detected by month 3.

The company also estimates that these techniques will also have a positive benefit on the maintenance component, but perhaps not so great as Case 2. They estimate a profit of 3MU.

The increased costs of the ICS are, in this case, 2MU. The cost of fix, should the error occur is 0.25MU. Thus, the overall profit (inclusive of the maintenance component of the contract) would be 9.8MU should the event occur and 10MU should it not.

Case 4

Alternatively, the company argues that it can dispense with the training course and use a more experienced person - a chief analyst/programmer - in exchange for one of the analyst programmers. The chief analyst/programmer costs 25% more than an analyst programmer, but is competent in the required techniques and is also experienced in providing successful on-the-job training. At worst the event again should be detected by month 3. The same benefit should apply to the maintenance phase as in Case 3.

The increased costs of the ICS are, in this case, 3MU. The cost of fix, should the error occur is 0.25MU. Thus, the overall profit (inclusive of the maintenance component of the contract) would be 8.8MU should the event occur and 9MU should it not.

In summary

Event occurs	Profit (MU)			
	ICS#1	ICS#2	ICS#3	ICS#4
Yes	(5.5)	6.8	9.8	8.8
No	9	8.3	10	9

Table 5: The bottom line effectiveness of the four candidate ICS (fixed price)

The company decides upon ICS#3.

An alternative scenario

It is interesting to consider what might happen if the contractual situation was quite different. What would happen, for example, if the contract was time and materials and there was no penalty clause. Suppose the company elects to charge its analyst programmers at a daily rate, equivalent to 1.67 MU per month, and its chief analyst programmers at a daily rate, equivalent to 2.09 MU per month. These rates allow for overhead and profit.

Ignoring maintenance, as that being on a time and materials basis, it would appear that the greater the number of bugs, the more maintenance work is required and therefore the greater the profit(!), we have:

ICS	Revenue	Cost	Profit
ICS#1 (no event)	60	51	9
ICS#1 (event occurs)	70.1	59.5	10.6
ICS#2 (no event)	66.8	56.7	10.1
ICS#2 (event occurs)	69.3	58.2	11.1
ICS#3 (no event)	60	53	7
ICS#3 (event occurs)	60.5	53.3	7.3
ICS#4 (no event)	65.1	54	11.1
ICS#4 (event occurs)	65.5	54.3	11.3

Table 6: The bottom line effectiveness of the four candidate ICS for the development phase (time and materials)

Note that in each case, the company makes a *greater* profit when the event occurs - it is almost as if they are being paid to do badly. ICS#3 is now the worst option, as it always results in the least profit. However, it does represent the least cost to the client.

With regard to the maintenance component of the contract, the company would charge 11.8 MU, resulting, in Case 1, with a profit of 1.8 MU. In Cases 3-4 less work is involved and therefore less profit (in fact 0.89, 1.25, 1.25MU respectively). In view of this, the company could elect to go fixed priced for Cases 3-4. The anticipated profit figures (and client charges including maintenance) would therefore be:

ICS	(event occurs)		(no event)	
	Profit	Client pays	Profit	Client pays
ICS#1	12.4	81.9	10.9	71.9
ICS#2	16.1	75.9	15.1	72.7
ICS#3	10.3	68.8	10.1	68.4
ICS#4	14.3	73.8	14.1	73.4

Table 7: The bottom line effectiveness of the four candidate ICS for a mix of time and materials (development phase) and fix price (maintenance). Note: ICS#1 is time and materials for both phases

Note that ICS#2 is the best from the perspective of making a handsome profit, whilst still not being the most expensive option from the perspective of the client. What we see here is the interplay between the client taking the risk during the development phase and the company taking the risk during the maintenance phase. The company is prepared to do this because of the superior ICSs involved in Cases 2-4. Note also that in a highly competitive situation, Case 3 wins as it allows the company to offer a low price whilst still demonstrating good value and competence.

Comment

Just who takes the risk in these situations is an important decision. For ICS#1 (see Table 7 above) the client takes all the risk. If the company (i.e. the contractor) makes an error, the client pays for it. With ICS#2-4, the company has a better ICS, which in Cases 2 and 4 does cost the client more. However, in all cases the client pays less than in Case 1 if the company does make an error. The question therefore is “*is it in the client’s best interest to pay the company to improve the effectiveness of its ICS?*” Table 7 suggests that the answer is not only “yes” but also that it can be done for next to nothing (compare ICS#3 with ICS#1).

Perhaps this example gives us an insight as to why customers do put pressure on large suppliers to introduce management systems (whether for risk, or subordinate areas such as quality or information security).

MEASURING IMPROVEMENT

Objective

If the current ICS is not achieving management’s objectives, management needs to be able to determine what needs to be changed and plot a course of action to implement those changes. As the ICS evolves, management again needs to measure the actual ICS to see how its effectiveness is improving.

Apart from the overall ICS, management may focus attention on particular controls. This may be in response to incidents or changes in threat.

Measurement

There are two types of improvement:

- An improvement to the overall ICS
- An improvement to an individual control.

Measurements are generally made using the methods previously described, i.e. using Table 2 to determine the actual class of a control, using Table 3 to determine operational effectiveness and Table 4 to determine cost-effectiveness. However, individual improvements may require additional metrics as discussed below.

What data do I need?

- Time to detect, time to fix, time window (either for individual events and/or averaged across many events of the same class)

- Cost of fix and impact penalty (again either for individual events and/or averaged across many events of the same class)
- Number and frequency of events
- Cost of control, cost of ICS
- Whether controls are protected by self-policing procedures and whether those procedures are fail-safe
- What incidents there have been that were not anticipated.

Improvement to the overall ICS

There are a variety of improvements that you may wish to make to the overall ICS and conform by measurement, e.g.

- Advancement to a higher Category
- Removal of redundant controls
- Increased cost effectiveness.

Advancement to a higher Category

In implementing such an improvement you will undoubtedly know what changes need to be made to effect the transition. Theoretically, you only need to re-measure the control class (Table 2) and re-apply the criteria (Table 3) to the changes. However, it is prudent to identify what criteria are borderline and monitor those as well. In that way, you can check that the operational effectiveness does not decrease as a result of changes made.

Removal of redundant controls

A control may be redundant if:

- The events that the control is designed to trap are universally trapped by some other control.
- The control does not trap other events.

The latter condition ensures that controls are not removed because they are redundant with respect to one event but not some other. If a control is truly redundant its removal should lead to improvements in cost effectiveness (which can be gauged using Table 4) together with similar improvements in business efficiency.

The first step would be to identify the controls that are candidates for removal and then monitor/measure how they work. Can you establish, for example, the number and frequency of the events that each control does trap? The second step is to establish a back-out plan such that, if removal of a control harbingers disaster, the control can be speedily put back! The third step is to remove the control and monitor/ measure how the other controls behave to establish confidence that the removal achieves its objectives.

Increased cost effectiveness

The idea here is to measure the cost effectiveness of either the whole ICS or a selected subset of it, by first acquiring the data necessary to apply Table 4. If improvements are required, the changes to particular controls are then identified. These are discussed below. On making the changes the measurements are repeated to confirm that the changes have met their objectives. Note that the removal of redundant controls is likely to have a positive impact on the overall cost effectiveness of the ICS.

Improvement to an individual control

There are a variety of improvements that you may wish to make to an individual control, or group of controls, and conform by measurement, e.g.

- Advancement to a higher Class or within Class
- Adding/changing self-policing and fail-safe properties to groups of controls
- Increased cost effectiveness.

Advancement to a higher Class or within Class

Moving from reactive to detective/ preventive, or moving from detective to preventive, is generally implemented by removing one control and replacing it by another. On the other hand, moving from one reactive class to another reactive class, or moving from one detective class to another detective class, is generally implemented merely by improving the particular characteristics of the control rather than changing it beyond recognition.

For example, moving from Class 7 to Class 6 may mean documenting what we did last time (or should have done in hindsight). Moving from Class 6 to Class 5 requires pre-deployment of some parts of the plan. The plan, however, is essentially the same in each case.

Likewise, moving from Class 4 to Class 3 and hence to Class 2 may be accomplished by improving the time to detect and/or the time to fix though educating and training staff, and essentially little change to anything else.

In each of the above cases we need only to measure the time parameters to indicate improvement or not as case be.

Self-policing and fail-safe properties

The value of a self policing procedure is that it promptly detects a control failure, allowing some other control to take over. As shown in Figure 6, if a control is not protected by a self-policing procedure, a control failure may go undetected and ultimately manifest itself following expiry of the time window. Is there evidence of this happening in the ICS? If there is, you not only have a need for a self-policing procedure but a means to monitor its effect. Once implemented the number of events trapped as Class 7 should reduce, ideally to zero. Note that self-policing and fail-safe properties are requirements of the higher order categories of ICS.

Increased cost effectiveness

Advancement to a higher Class or improvements within class are likely to have a beneficial effect on the cost-effectiveness of the control. Other improvements may be made by paying attention to the costs involved, in particular the cost of the control. Note that reducing the time to detect (see our cost-effectiveness worked example, page 18) may automatically result in a shorter time to correct and a lower cost of fix. Use Table 4.

A Worked Example (continued again)

Let us return to our software company example and this time consider the candidate ICS identified when we were considering the cost-effectiveness problem. The category for ICS#1 is A, as determined previously. The chosen ICS (#3) adds a new control, which is Class 2 non-degradable. By itself this is insufficient to change the category of the ICS. However, it was chosen because

it ought to increase the cost effectiveness of the ICS. We need to monitor the cost of the control, the cost of fix and any impact penalty to confirm this. We can also monitor the improvements in the control itself by recording the time to detect and the time to fix for the events that it discovers.

Other questions that we ought to monitor are:

- Are there errors that it was designed to detect that escape detection?
- Are these trapped by the software testing control?
- What proportion of these escape detection and are ultimately found by the customer?

RISK TREATMENT PLANS

In this section we propose a methodology for generating Risk Treatment Plans, which makes use of the fundamental theory which we have just discussed.

We have used the methodology extensively in the information security arena. We have taught senior managers/risk owners how to use it in various parts of the world and they have been able to apply it, not only in the context of information security but also to other business/governance concerns. We therefore believe that this methodology *works* and is *teachable*, *repeatable* and *reproducible*.

We start by saying a little more about events and impacts on which the methodology is founded.

Events

The events referred to in this paper are bad things that cause trouble. The insert (below) lists those events, which in our BS7799 work we feel are common across many businesses. In addition we would add other events that were specific to that particular organisation. An example would be "*one of our aircraft has broken down in the Indian Ocean*". We told a story about this earlier.

Events that are likely to be common across many businesses are:

- Theft
- Acts of God, vandals and terrorists
- Regular fraud
- IT failure
- Hacking
- Denial of Service attacks
- Disclosure
- Breach of the law

Typically, any occurrence of such events would be reported to management, the speed of reporting being a function of their severity. Think of the event as a newspaper headline.

Impacts

Likewise, it is possible to characterise the damage, or impact of an event in a standard manner. The insert (below) lists those impacts, which again in our BS7799 work we feel are common across many businesses.

Impacts that are likely to be common across many businesses are:

- Customer dissatisfaction
- Adverse press coverage
- Loss of revenue
- Unanticipated costs
- Inability to carry out some or all of its business
- Loss of the monetary value of buildings and contents
- Failure to prosecute
- Court action against an employee or the business itself

The occurrence of an event may give rise to several impacts and may also trigger other events.

Risk Treatment Plans

Risk treatment is an ISO term that means the “*treatment process of selection and implementation of measures to modify risk* [ISO Guide 73]”. We can use this concept to develop a simple methodology for applying our fundamental theory. Figure 10 shows a fragment from our stylised form of a BS7799-2:2002 Risk Treatment Plan (RTP).




Figure 10: A fragment of a stylised Risk Treatment Plan

The process of producing the RTPs can be described in terms of a series of steps.

Step 1 - identify the events

Name the event and briefly describe it. Our usual approach is to start with the standard events described above and augment them with client specific concerns.

Step 2 - identify the assets

We usually start with a generic list that includes such things as:

- Buildings and contents
- IT hardware and networks
- Infrastructure and application software
- Computerised data concerning the organisation’s business
- Paper documents and records concerning the organisation’s business
- Supporting data, documentation and records.

We will add to this list and otherwise modify it as necessary, the idea ultimately being to derive the assets that require protection from the analysis, rather than the other way round (which unfortunately seems to be the conventional way of carrying out a risk assessment).

Step 3 - identify the impacts

Our usual approach is to start with the standard impacts described above and augment them with client specific impacts as required.

Step 4 - identify the threats

We usually start with a generic list of threat agents that includes such entries as:

- Fire, flood and other forms of “natural disaster”
- Power and other utility failure
- Customers and suppliers
- Disaffected staff
- Spies
- Thieves
- Vandals and terrorists
- Hackers
- Errors and mistakes.

We will add to this list and otherwise modify it as necessary.

Step 5 - produce the RTPs

This step is repeated for each event.

First (see Figure 10), write down the description of the event and list the assets that are affected. Augment/modify the asset inventory if there is an asset that we wish to refer to that is not already in the list.

Second, document the applicable impacts and order them in the priority they are to receive. Record if any are to receive equal priority treatment.

Third, list the applicable threats.

Fourth, repeat the steps 5a - 5d below until all the impacts have been dealt with. If the impacts are listed in priority order, take them in that order. If two or more have the same priority, take them together.

Note also that:

- In practice the analysis for a given event-impact pair will break down into a number of “threads”. Each thread takes some particular starting point to its logical conclusion, which is a statement of acceptable or unacceptable residual risk.
- A useful starting point (which we illustrate in Step 5a below) is to consider some way in which a given threat agent can cause, in the context of the event, the impact under consideration. Once this thread has been considered, the next thread would consider some other way in which the threat agent can cause the same impact. Finally, we should include a thread for which the threat agent is unknown.
- Quite often RTPs become repetitive. We obviate this as much as possible through the order in which we deal with the events and through cross references between RTPs, where the procedures are the same albeit dealing with another thread or event/impact pair.

Step 5a - identify the risks leading to a particular impact (or impacts if the impacts have the same priority) for known threats

Consider the event and the impact(s).

The first question to ask is “what is being done about it already?”.

Unless we are really starting from a clean sheet of paper (which we might be doing in the case of a new system), there will already be procedures and technology in place to deal with the event-impact. Even if we are considering something new, there may be policies that we are obliged to follow that dictate what procedures and technology we must put in place. If not, then we are free to identify what we need. In all three cases we will assume for the purpose of describing this methodology that those procedures/ technology are (or will be) documented.

We then ask, what do these procedures and technology accomplish in terms of this particular event-impact? and write down the answer. We find it best to tell it like a story. Say how a threat agent, in the context of this event, could bring about the impact under consideration, e.g. “*A hacker could bring about the inability of the organisation to carry out some or all of its business by mounting a denial of service (DoS) attack on the network.*” You then write down, in as few words as possible, what the (established) procedures and technology will do about this, e.g. “*The first line of defence against such an attack is the firewall.*” Note that this is a Class 1 control.

We then document the risk, e.g. “*Our ISP provides this firewall as a managed service. We do not therefore know whether this firewall is always correctly configured, or if is under attack.*” Note that this is tantamount to saying “what if the first control does not work?”.

We then ask if this risk is acceptable or not. If it is acceptable, say so and say why, e.g. “*Nevertheless, this lack of knowledge is considered to be an acceptable risk because there is a second line of defence, which lies in hardening the network components in accordance with the IT policy for “Hotfix and service pack upgrades”.* Note in this example the introduction of another control to address the failure of the first.

The analysis proceeds in this way until all of the controls that are used (or are to be used) and their effects have been documented, e.g. following on from the previous example the RTP might say “*However (a) Currently the internal network IP addresses are public addresses. This presents an unacceptable risk and therefore we need to convert these addresses to private addresses; (b) There is a risk that a hacker could still exploit some known vulnerability for which the Hotfix had not yet been applied or exploit some other yet unreported vulnerability for which there is currently no Hotfix. At present, this is an acceptable risk because of the low profile of our web site and those that it hosts for our customers.*”

Note that this “thread” terminates with both an unacceptable residual risk and an acceptable residual risk.

The next thread might consider hostile code insertion.

Step 5b - identify the risks leading to a particular impact for unknown threats

It is prudent to ensure that appropriate controls are in place to deal with the situation where the event has occurred for some unanticipated reason, i.e., the threat agent and/or the attack method was not known or anticipated at the time the analysis is performed. It is possible that these will have already been identified during Step 5a, in which case Step 5b merely identifies what they are and explains how their workings are independent of threat agent and/or attack method. If the controls identified in the Step 5a threads are not suitable, check whether others

exist in the ICS that are suitable and document them. If none can be found decide whether the residual risk is acceptable or not.

Step 5c - dealing with unacceptable residual risks

If a thread terminates with an unacceptable residual risk we need to do something about it. We usually approach this by having a “To-Do-List”. We decide what needs to be done about the unacceptable risk - which at the very least will be to investigate the options - and append it to the To-Do-List. It is then a question of project managing the To-Do-List. Of course, whilst a particular problem is being resolved, we are running an unacceptable risk. It is possible that we cannot do anything about that apart from keeping our fingers crossed. Otherwise, it would be appropriate to introduce some short-term measure.

Step 5d - optimising the ICS

The RTP thus far describes the control structure that exists and, via the To-Do-List, also that which is planned for the future.

Identify the class (1-7) for each control. Sometimes we also find it useful to draw the Venn diagrams to show how the controls interact along a given thread.

Use the event and impact data and Table 4 to determine whether a different class for a particular control would be more appropriate. Make your decisions in the context of the other controls, in particular the other controls in the same thread.

If we are dealing with a real life situation, the decision to change the control structure is recorded by making the appropriate entries in the To-Do-List. We refer to such changes as “improvements”.

If we are dealing with a future system (for example, if we are in the process of working out the ICS requirements for a new IT system) then we would merely change the control structure and iterate steps 5a - 5d.

Step 6 - tidy up

Once all the risk treatment plans have been developed, there may be a certain amount of tidying up to do.

First, check that all the assets in the asset inventory have been used. If any are left over, ensure that that is not because of some oversight and, if not, remove them from the inventory. Note that new ones, not present in the initial list, may have been defined on the fly. Ensure that all the RTPs that ought to refer to these additional assets do so. Note that in this way we use the original list merely as a starting position. Following augmentation and tidy-up we effectively finish with a list that is derived from the risk assessment. In other words we identify those assets that require protection in order to ensure acceptability of risk rather than assume what needs protection and use the risk assessment to identify how they should be protected.

Second, check that all the impacts in the impact list have been dealt with. If there are any discrepancies, or additional impacts have been added on the fly, proceed as in the case of tidying up the assets.

Third, check that all the threat agents in the threat list have been dealt with. If there are any discrepancies, or additional impacts have been added on the fly, proceed as in the case of tidying up the assets and impacts.

Fourth, ensure that all event-impact pairs have been dealt with.

Fifth, check that all threads end in a statement of acceptable/unacceptable risk.

Sixth, check that all control failures have been considered.

Finally, check that the To-Do-List entries remedy all unacceptable residual risks and implement all identified improvements.

CONCLUSIONS

This paper has set out:

- the methodology to measure the effectiveness of the control element of an ICS
- a methodology to present risk options in the form of a story as a RTP to improve communications between risk specialists and senior management.

Internal Control system

An ICS is a mandatory requirement to meet the obligations of Corporate Governance and the legislation, throughout the world, requiring Directors and Senior Managers to maintain effective control of the organisation and to demonstrate positively their involvement in the control of the organisation.

The ICS can have a material impact on the ability of an organisation to meet its objectives. The paper shows that you can have an ICS that inhibits an organisation to meet its objectives as well as an ICS that assists. Almost certainly all organisations will need to be able to react promptly to unexpected events.

ICS Metrics

We propose two sets of metrics for use in determining the effectiveness of an ISMS within an organisation. The first set of metrics is independent of external factors and is therefore a true measure of the effectiveness of the organisation's procedures and management system.

“Operational effectiveness is determined solely by measuring the time parameters.”

These metrics are Time dependent. They are the time to detect an event and the time taken then to rectify the consequences. We anticipate that analysis against these metrics will be by class of event. This led us to see clearly the view expressed on empirical evidence that prompt detection of potential events is the best solution and the optimum position is that a procedure is constructed so that any errors made are automatically detected (for example: the old fashioned double entry bookkeeping system). They are useful in designing an appropriate ICS and verifying that the implementation accords with the design.

The second set concerns costs and impact penalties and is useful in deciding whether the ICS is cost effective.

Classes of ICS

We knew that all business operations incur a cost. We divided the cost into four categories:

- Cost of doing business
- Cost of having an ICS
- Impact penalty from a control failure which will materialise if the event is not detected within the time window where rectification is possible

- Cost of rectification following the manifestation of an event.

We postulate seven classes of controls with differing properties. These range from a procedure to immediately detect an event and stop the event impacting the organisation (clearly an optimum position) to a catastrophe situation when the event may cause a business failure.

This categorisation enables people to consider the nature of ICS procedures necessary to address identified events in order to optimise the cost effectiveness of the control procedures. It must always be remembered that some events will be outside the control of the organisation, therefore prevention is not an option and the ICS procedures are forced to address rectification only (if this is possible at all).

Risk Treatment Plans

The use of Risk Treatment Plans (RTPs) expressed as a story of what organisation has put in place to address risk events (of their choosing) in relation to the possible impacts of that event allows everyone from the Board downward to understand what risk management issues are addressed and how. The granularity of the RTPs is a matter that an organisation (or part thereof) may elect to meet their needs. The story will always address the questions ‘Suppose the control does not work? What do we do then? This enables a thorough systematic decomposition of the procedures in place (or proposed) so as to enable people to confirm that the controls are just sufficient for the purpose.

Applicability

We believe that this methodology is applicable to all risk situations whether these be the risks of doing business, information security, quality, environmental, legislative/regulatory compliance etc.

We trust that this paper will make a substantial contribution to the ongoing debate on how Information Security (or Assurance) can be achieved by organisations in the future.

Acknowledgments

We would like to record our thanks to the very many people, all over the world, who have over the years taught us about Information Security and Internal Control structures.

In particular in respect of this paper we would record our thanks to

Matthew Pemble of RBS

Richard Hackworth of HSBC

David Spinks of EDS

Harvey Mattinson of the UK Cabinet Office

Michael Nash of Gamma Secure Systems

Leslie McCartney of Reuters

Ted Humphries of Xisec

Marc Kekicheff of Visa International

Regina Brewer of Konami (Europe)

We trust that our contribution will be seen as our thanks to them for their training to us.

References

- [APB guidance] Briefing paper - Providing Assurance on the effectiveness of Internal Control issued by the Audit Practices Board July 2001, see <http://www.apb.org.uk/> Copies from ABG Professional Information info@abgpublications.co.uk
- [BASEL 2] The Bank of International Settlements, the New Basel 2 Accord, see <http://www.bis.org/>
- [BS7799-2] Information security management systems - Specification with guidance for use, BS 7799-2:2002
- [Higgs] Review of the role and effectiveness of non-executive directors, Derek Higgs, 2003, see <http://www.dti.gov.uk/>
- [GlobalPlatform] The GlobalPlatform Card Specification and the GlobalPlatform Card Security Security Requirements Specification, <http://www.globalplatform.org/>
- [ISO 14001] Environmental management systems - Specification with guidance for use, BS EN ISO 14001:1966
- [ISO/IEC 17799] Information technology - Code of practice for information security management, BS ISO/IEC 17799:2000
- [ISO 9001] Quality management systems - Requirements, BS EN ISO 9001:2000
- [OECD] Organisation for Economic Co-operation and Development, Corporate Governance, see <http://www.oecd.org/>
- [Sarbanes-Oxley] Sarbanes-Oxley Act of 2002, USA Congress, an Act to protect investors by improving the accuracy and reliability of corporate disclosures made pursuant to the securities laws, and for other purposes, see <http://news.findlaw.com/>
- [Turnbull] Internal Control, Guidance for directors on the Combined Code (The Turnbull Report), Institute of Chartered Accountants in England and Wales, see <http://www.icaew.co.uk/>

About the authors

Dr. David Brewer is a founder director of Gamma. He has been involved in information security since he left university, and is an internationally recognised consultant in that subject. He was part of the team who created the ITSEC and the Common Criteria, and has worked for a wide range of government departments and commercial organisations both at home and abroad.



Mr. William List, CA hon FBCS, is the proprietor of Wm. List & Co. He has been involved in security and audit for some 40 years. He has been involved in the development of secure business applications and the development of various accounting and IT standards. He retired as a partner in KPMG. He is a previous chairman of the BCS security expert panel. Both are currently part of the international team developing the 7799 family of standards, and are two of the driving forces behind the Part 2 ISMS standard. They provided training in implementing ISO/IEC 17799 and have assisted many clients to build ISMSs since 1998 in Europe, East Africa and the Far East.

Dr Brewer can be contacted at: dbrewer@gammassal.co.uk, www.gammassal.co.uk

Mr list can be contacted at: w.list@ntlworld.com



◆ A SPECIALIST GROUP OF THE BCS ◆

Membership Application

(Membership runs from July to the following June each year)

I wish to APPLY FOR membership of the Group in the following category and enclose the appropriate subscription.

CORPORATE MEMBERSHIP (Up to 5 members)* £75

*Corporate members may nominate up to 4 additional recipients for direct mailing of the Journal (see over)

INDIVIDUAL MEMBERSHIP (NOT a member of the BCS) £25

INDIVIDUAL MEMBERSHIP (A members of the BCS) £15

BCS membership number: _____

STUDENT MEMBERSHIP (Full-time only and must be supported by a letter from the educational establishment).

Educational Establishment: _____ £10

Please circle the appropriate subscription amount and complete the details below.

INDIVIDUAL NAME: (Title/Initials/Surname)
POSITION:
ORGANISATION:
ADDRESS:
POST CODE:
TELEPHONE: (STD Code/Number/Extension)
E-mail:
PROFESSIONAL CATEGORY: (Please circle) 1 = Internal Audit 4 = Academic 2 = External Audit 5 = Full-Time Student 3 = Data Processor 6 = Other (please specify)
SIGNATURE: _____ DATE: _____

PLEASE MAKE CHEQUES PAYABLE TO "BCS IRMA" AND RETURN WITH THIS FORM TO

Janet Cardell-Williams, IRMA Administrator, 49 Grangewood, Potters Bar, Herts EN6 1SL. Fax: 01707 646275

ADDITIONAL CORPORATE MEMBERS

INDIVIDUAL NAME: (Title/Initials/Surname)
POSITION:
ORGANISATION:
ADDRESS:
POST CODE:
TELEPHONE: (STD Code/Number/Extension)
E-mail:
PROFESSIONAL CATEGORY: 1 = Internal Audit 4 = Academic 2 = External Audit 5 = Full-Time Student 3 = Data Processor 6 = Other (please specify)

INDIVIDUAL NAME: (Title/Initials/Surname)
POSITION:
ORGANISATION:
ADDRESS:
POST CODE:
TELEPHONE: (STD Code/Number/Extension)
E-mail:
PROFESSIONAL CATEGORY: 1 = Internal Audit 4 = Academic 2 = External Audit 5 = Full-Time Student 3 = Data Processor 6 = Other (please specify)

INDIVIDUAL NAME: (Title/Initials/Surname)
POSITION:
ORGANISATION:
ADDRESS:
POST CODE:
TELEPHONE: (STD Code/Number/Extension)
E-mail:
PROFESSIONAL CATEGORY: 1 = Internal Audit 4 = Academic 2 = External Audit 5 = Full-Time Student 3 = Data Processor 6 = Other (please specify)

INDIVIDUAL NAME: (Title/Initials/Surname)
POSITION:
ORGANISATION:
ADDRESS:
POST CODE:
TELEPHONE: (STD Code/Number/Extension)
E-mail:
PROFESSIONAL CATEGORY: 1 = Internal Audit 4 = Academic 2 = External Audit 5 = Full-Time Student 3 = Data Processor 6 = Other (please specify)



◆ A SPECIALIST GROUP OF THE BCS ◆



Management Committee

CHAIRMAN	Alex Brewer	alex.brewer@morganstanley.com
SECRETARY	Siobhan Tracey	siobhan.tracey@booker.co.uk
TREASURER	Jean Morgan	jean@wilhen.co.uk
MEMBERSHIP	Celeste Rush	rushlse97@aol.com
JOURNAL EDITOR & SECURITY PANEL LIAISON	John Mitchell	john@lhscontrol.com
WEBMASTER	Allan Boardman	allan@internetworking4u.co.uk
EVENTS PROGRAMME CONSULTANT	Raghu Iyer	raguriyer@aol.com
LIAISON - IIA & NHS	Mark Smith	mark.smith@lhp.nhs.uk
LIAISON - LOCAL AUTHORITY	Peter Murray	cass@peterm.demon.co.uk
LIAISON - ISACA	Ross Palmer	ross.palmer@hrplc.co.uk
MARKETING	Wal Robertson	williamr@bdq.com
ACADEMIC RELATIONS	David Chadwick	d.r.chadwick@greenwich.ac.uk
	David Lilburn Watson	dlwatson@bcm.co.uk
SUPPORT SERVICES		
ADMINISTRATION	Janet Cardell-Williams t: 01707 852384 f: 01707 646275	admin@bcs-irma.org
OR VISIT OUR WEBSITE AT	www.bcs-irma.org	Members' area Userid = irmamembers Password = irma2004

BCS IRMA SPECIALIST GROUP ADVERTISING RATES

Reach the top professionals in the field of EDP Audit, Control and Security by advertising in the BCS IRMA SG Journal. Our advertising policy allows advertising for any security and control related products, service or jobs.

For more information, contact John Mitchell on 01707 851454, fax 01707 851455 email john@lhscontrol.com.

There are three ways of advertising with the BCS IRMA Specialist Group:

The Journal is the Group's award winning quarterly magazine with a very defined target audience of 350 information systems audit, risk management and security professionals.

Display Advertisements (Monochrome Only) Rates:

- Inside Front Cover £400
- Inside Back Cover £400
- Full Page £350 (£375 for right facing page)
- Half page £200 (£225 for right facing page)
- Quarter Page £125 (£150 for right facing page)
- Layout & artwork charged @ £30 per hour

Inserts can be included with the Journal for varying advertising purposes, for example: job vacancies, new products, software.

Insertion Rates:

For inserts weighing less than 60grams a flat fee of £300 will be charged. Weight in excess of this will incur additional charges:

- 60-100grams: 14p per insert
- 101-150g: 25p per insert
- 151-300g: 60p per insert
- 301-400g 85p per insert
- 401-500 105p per insert

Thus for an insert weighing 250g it would cost the standard £300 plus weight supplement of £210 (350 x 60pence) totalling £510.

Discounts:

Orders for Insert distribution in four or more consecutive editions of the Journal, if accompanied by advance payment, will attract a 25% discount on quoted prices.

Direct mailing

We can undertake direct mailing to our members on your behalf at any time outside our normal distribution timetable as a 'special mailing'. Items for distribution **MUST** be received at the office at least 5 WORKING DAYS before the distribution is required. Prices are based upon an access charge to our members plus a handling charge. Access Charge £350. Please note photocopies will be charged at 21p per A4 side.

Personalised letters:

We can provide a service to personalise letters sent to our members on your behalf. This service can only be provided for standard A4 letters, (i.e. we cannot personalise calendars, pens etc.). The headed stationery that you wish us to use must be received at the Office at least ten working days before the distribution is required. Please confirm quantities with our advertising manager before dispatch. If you require this service please add £315 to the Direct mailing rates quoted above.

Discounts: Orders for six or more direct mailings will attract a discount of 25% on the quoted rates if accompanied by advance payment

Contacts

Administration

Janet Cardell-Williams,
49 Grangewood, Potters Bar, Hertfordshire EN6 1SL
Email: admin@bcs-irma.org
Website : www.bcs-irma.org

BCS IRMA Specialist Group Advertising Manager

Eva Nash Tel: 01707 852384
Email: admin@bcs-irma.org